arXiv:1604.06741v1 [math.PR] 22 Apr 2016

The Incipient Giant Component in Bond Percolation on General Finite Weighted Graphs

David J. Aldous *

April 25, 2016

Abstract

On a large finite connected graph let edges e become "open" at independent random Exponential times of arbitrary rates w_e. Under minimal assumptions, the time at which a giant component starts to emerge is weakly concentrated around its mean.

MSC 2010 subject classifications: 60K35, 05C80.

1 Introduction

Take a finite connected graph $(\mathbf{V}, \mathbf{E})$ with edge-weights $\mathbf{w} = (w_e)$, where $w_e > 0 \ \forall e \in \mathbf{E}$. To the edges $e \in \mathbf{E}$ attach independent Exponential(rate w_e) random variables ξ_e. In the language of percolation theory, say that edge e becomes *open* at time ξ_e. The set of open edges at time t determines a random partition of $\mathbf{V}$ into connected components; write $C(t)$ for the largest number of vertices in any such connected component. Now consider a sequence $(\mathbf{V}_n, \mathbf{E}_n)$ of such weighted graphs, where both the graph topologies and the edge-weights are arbitrary subject only to the conditions that $|\mathbf{V}_n| \to \infty$ and that for some $0 < t_1 < t_2 < \infty$

$$\lim_n \mathbb{E}C_n(t_1)/|\mathbf{V}_n| = 0; \quad \bar{c} := \liminf_n \mathbb{E}C_n(t_2)/|\mathbf{V}_n| > 0. \tag{1}$$

In the language of random graph theory, this condition says that a *giant component* emerges (with non-vanishing probability) sometime between t_1 and t_2. Proposition 1 asserts, informally, that the "incipient" time at which the giant component starts to emerge is deterministic to first order.

*Department of Statistics, 367 Evans Hall # 3860, U.C. Berkeley CA 94720; aldous@stat.berkeley.edu; www.stat.berkeley.edu/users/aldous. Aldous's research supported by N.S.F Grant DMS-1504802.

Proposition 1 *Given a sequence of graphs satisfying (1), there exists a deterministic sequence* $\tau_n \in [t_1, t_2]$ *such that, for every sequence* $\omega_n \uparrow \infty$ *sufficiently slowly, the random times*

$$T_n := \inf\{t : C_n(t) \geq |\mathbf{V}_n|/\omega_n\}$$

satisfy

$$T_n - \tau_n \to_p 0.$$

In the special cases of the complete graph and the 2-dimensional discrete torus (with constant edge-weights) we are essentially dealing with component sizes in the classical Erdős-Rényi and the bond percolation on $\mathbb{Z}^2$ processes, for which much stronger results are known about the "scaling window" of time over which the giant component emerges [10, 11]. Such stronger results have been generalized (again with constant edge-weights) in several directions, for instance to random subgraphs of certain transitive finite graphs [12, 13] or to random subgraphs of graphs under assumptions that force the critical subgraphs to be "tree-like" as in the Erdős-Rényi case [14]. Proposition 1 gives a comparatively weak concentration property which one would expect to hold in every "natural" example, but it is perhaps remarkable that it holds in the generality stated. As a comparison, in the classical cases one also has the same "weak concentration" property for the random times

$$T_n^*(s) := \inf\{t : C_n(t) \geq s|\mathbf{V}_n|\}$$

for fixed $0 < s < 1$. One would expect this to extend to other "natural" examples, but a simple example outlined in section 3.2 shows it does not hold in the generality of Proposition 1, even if we assume the weighted graph to be vertex-transitive. Some conjectures concerning the post-incipient regime are given in section 3.4.

The proof is based on a simple general variance bound described in section 2.1. A technically more complicated application of that bound to first passage percolation on general weighted graphs, plus other simple applications, can be found in [2]. "Big picture" discussions of various random processes over finite edge-weighted graphs can be found in [1] and [3].

2 Proof of Proposition 1

We divide the proof into three steps.

2.1 Step 1: The general variance bound for increasing set-valued processes.

Fix a weighted graph $(\mathbf{V}, \mathbf{E}, \mathbf{w})$ and $1 < \omega < |\mathbf{V}|$. The process of open edges in our bond percolation process is a continuous-time Markov chain, Z_t say, whose state space is the set of subsets $S \subseteq \mathbf{E}$ and whose transition rates are

$$S \to S \cup \{e\} : \quad \text{rate } w_e, \quad (e \notin S).$$

We seek to study the distribution of the stopping time

$$T = \inf\{t : C(t) \geq |\mathbf{V}|/\omega\} \tag{2}$$

when this chain starts in state $\emptyset$. It makes sense to also consider this mean hitting time started from an arbitrary subset S of open edges, that is

$$h(S) := \mathbb{E}_S T$$

which clearly has the property

$$h(S') \leq h(S) \text{ whenever } S \to S' \text{ is a possible transition.} \tag{3}$$

There is a general concentration inequality for Markov chain stopping times with this property: it is an easy consequence of martingale identities for $\mathbb{E}T$ and var T, though apparently not well known. In general the starting state is arbitrary, but to fit our setting we take it as $\emptyset$.

Lemma 2 ([2] equation (9) and subsequent display) *Given property (3), for arbitrary $\delta > 0$,*

$$\frac{\text{var}_\emptyset\ T}{(\mathbb{E}_\emptyset T)^2} \leq \delta + \frac{\mathbb{E}_\emptyset \int_0^T q_\delta(Z_u) du}{\mathbb{E}_\emptyset T}$$

where

$$q_\delta(S) := \sum_{S' :\ h(S) - h(S') > \delta \mathbb{E}_\emptyset T} q(S, S')(h(S) - h(S')) \leq 1 \tag{4}$$

and $q(S, S')$ are the transition rates.

We will use a general consequence, derived as Corollary 3 below, Note that $h(S) \leq h(\emptyset) = \mathbb{E}_\emptyset T$ for subsets S under consideration. Now consider the first time (if ever) W_δ that the process makes some transition $S \to S'$ with $h(S) - h(S')$ greater than $\delta \mathbb{E}_\emptyset T$:

$$W_\delta := \inf\{t : h(Z_{t-}) - h(Z_t) > \delta \mathbb{E}_\emptyset T\}. \tag{5}$$

Then

$$\int_0^T q_\delta(Z_u)du = \int_0^{T\wedge W_\delta} q_\delta(Z_u)du + \int_{T\wedge W_\delta}^T q_\delta(Z_u)du$$
$$\leq \int_0^{T\wedge W_\delta} q_\delta(Z_u)du + T\mathbb{1}_{\{W_\delta<T\}}.$$

Property (3) implies the distribution of T has the submultiplicativity property

$$\mathbb{P}_\emptyset(T > t_1 + t_2) \leq \mathbb{P}_\emptyset(T > t_1)\ \mathbb{P}_\emptyset(T > t_2), \quad t_1, t_2 > 0 \tag{6}$$

and it is straightforward to show there is a certain "universal" function $\gamma(u) \downarrow 0$ as $u \downarrow 0$ such that, for any submultiplicative T and any event A we have $\mathbb{E}[T\mathbb{1}_A] \leq \gamma(\mathbb{P}(A))\ \mathbb{E}T$. So in our setting

$$\mathbb{E}_\emptyset[T\mathbb{1}_{\{W_\delta<T\}}] \leq \gamma(\mathbb{P}_\emptyset(W_\delta < T))\ \mathbb{E}_\emptyset T.$$

Then from Lemma 2

$$\frac{\mathrm{var}_\emptyset\ T}{(\mathbb{E}_\emptyset T)^2} \leq \delta + \gamma(\mathbb{P}_\emptyset(W_\delta < T)) + \frac{\mathbb{E}_\emptyset \int_0^{T\wedge W_\delta} q_\delta(Z_u)du}{\mathbb{E}_\emptyset T}.$$

Now consider

$$\widetilde{q}_\delta(S) := \sum_{S':\ h(S)-h(S')>\delta\mathbb{E}_\emptyset T} q(S, S')$$

so that

$$q_\delta(S) \leq \widetilde{q}_\delta(S)\ \mathbb{E}_\emptyset T.$$

But $\widetilde{q}_\delta(Z_u)$ is the intensity rate of W_δ, that is

$$\mathbb{1}_{\{W_\delta\leq t\}} - \int_0^{W_\delta\wedge t} \widetilde{q}_\delta(Z_u)du \text{ is a martingale,}$$

and so from the optional sampling theorem

$$\mathbb{E}_\emptyset \int_0^{W_\delta\wedge T} \widetilde{q}_\delta(Z_u)du = \mathbb{P}_\emptyset(W_\delta \leq T).$$

Combining this with the previous two displayed inequalities gives

Corollary 3

$$\frac{\mathrm{var}_\emptyset\ T}{(\mathbb{E}_\emptyset T)^2} \leq \delta + \gamma(\mathbb{P}_\emptyset(W_\delta \leq T)) + \mathbb{P}_\emptyset(W_\delta \leq T). \tag{7}$$

We remark that the argument for inequality (7) does not use any structure of the bond percolation process except property (3).

2.2 Step 2: Bounding in terms of the growth rate of the incipient giant component.

We now use the structure of the bond percolation process by relating T defined at (2) to

$$T^{(2)} = \inf\{t : C(t) \geq 2|\mathbf{V}|/\omega\} \geq T.$$

Consider a possible transition $S' \to S'' = S' \cup \{e\}$. We will show

$$h(S') - h(S'') \leq \mathbb{E}_{S''}(T^{(2)} - T). \tag{8}$$

There is a natural coupling $(Z'_u, Z''_u, u \geq 0)$ of the processes started from S' and from S''; that is, $Z''_u = Z'_u \cup \{e\}$ until the Exponential(w_e) time at which $e \in Z'_u$, after which time $Z''_u = Z'_u$. Write $C'(u), C''(u)$ for the largest component sizes, and T', T'' for the stopping times (2), applied to these coupled processes. At time $T^{(2)\prime\prime} = \inf\{t : C''(t) \geq 2|\mathbf{V}|/\omega\}$ the process Z'' contains a component of size at least $2|\mathbf{V}|/\omega$, and so after deleting edge e there must remain a component of size at least $|\mathbf{V}|/\omega$ of Z'. That establishes the second inequality in

$$T'' \leq T' \leq T^{(2)\prime\prime}$$

and the first inequality is immediate. Now

$$h(S') - h(S'') = \mathbb{E}T' - \mathbb{E}T'' \leq \mathbb{E}T^{(2)\prime\prime} - \mathbb{E}T'' = \mathbb{E}_{S''}(T^{(2)} - T)$$

establishing (8).

Now let V be the time (if any) that the bond percolation process started at $\emptyset$ makes a specified transition. Then (8) says that

$$h(Z_{V-}) - h(Z_V) \leq \mathbb{E}_\emptyset(T^{(2)} - T|\mathcal{F}_V) \text{ on } \{V \leq T\}. \tag{9}$$

Now fix $\delta > 0$. For each pair $\pi = (S', S' \cup \{e\})$ for which $h(S') - h(S' \cup \{e\}) \geq \delta \mathbb{E}_\emptyset T$ there is a time V_π as above. The random variable W_δ at (5) is $W_\delta = \min_\pi V_\pi$, and so (9) implies

$$\delta \mathbb{E}_\emptyset T \leq h(Z_{W_\delta -}) - h(Z_{W_\delta}) \leq \mathbb{E}_\emptyset(T^{(2)} - T|\mathcal{F}_{W_\delta}) \text{ on } \{W_\delta \leq T\}. \tag{10}$$

This in turn implies

$$\mathbb{P}_\emptyset(W_\delta \leq T) \leq \frac{\mathbb{E}_\emptyset(T^{(2)} - T)}{\delta \mathbb{E}_\emptyset T}.$$

Applying (7), and setting

$$\gamma^*(u) = \gamma(u) + u \downarrow 0 \text{ as } u \downarrow 0,$$

we find

$$\frac{\mathrm{var}_\emptyset\ T}{(\mathbb{E}_\emptyset T)^2} \le \delta + \gamma^* \left(\frac{\mathbb{E}_\emptyset(T^{(2)} - T)}{\delta \mathbb{E}_\emptyset T} \right)$$

Because δ is arbitrary, this implies

$$\frac{\mathrm{var}_\emptyset\ T}{(\mathbb{E}_\emptyset T)^2} \le \Gamma \left(\frac{\mathbb{E}_\emptyset(T^{(2)} - T)}{\mathbb{E}_\emptyset T} \right) \tag{11}$$

where

$$\Gamma(x) := \inf_{\delta > 0} (\delta + \gamma^*(x/\delta)) \downarrow 0 \text{ as } x \downarrow 0.$$

2.3 Step 3: A compactness reduction.

The remainder of the proof uses only "soft" arguments. We are given a sequence of weighted graphs satisfying (1). To emphasize dependence on ω_n write

$$T_n(\omega_n) := \inf\{t : C_n(t) \ge |\mathbf{V}_n|/\omega_n\}.$$

Take $\omega_n \ge 2$ to avoid trivialities. By the second condition in (1) and submultiplicativity (6) there is an integrable T^* such that

$$T^* \text{ stochastically dominates } T_n(\omega_n), \text{ for all } n,\ \omega_n. \tag{12}$$

By the first condition in (1) we can take $\omega_n \uparrow 0$ sufficiently slowly that $\mathbb{P}(T_n \le t_1) \to 0$. Looking at (11), we see that the proof of Proposition 1 reduces to the proof of

(*) for all $\omega_n \uparrow \infty$ sufficiently slowly, $\mathbb{E}(T_n(\frac{1}{2}\omega_n) - T_n(\omega_n)) \to 0$.

Property (12) implies compactness with respect to weak convergence and convergence of expectations. By a standard compactness principle, to prove (*) it will suffice to prove that every subsequence has a further sub-subsequence in which (*) holds, and – up to a change in notation – it is enough to show that the original sequence has some subsequence in which (*) holds.

Consider the set of all possible subsequential weak limits of sequences $T_{m_n}(\omega_n)$. This is compact, so has at least one element μ which is maximal with respect to the "stochastic order" partial order. And a subsequence $T_{m_n}(\omega_n^*)$ converging to μ clearly has property (*), because for any $\omega_n \le \omega_n^*$ we have $T_{m_n}(\omega_n) \ge T_{m_n}(\omega_n^*)$ and so by maximality $T_{m_n}(\omega_n)$ must also converge to μ, as must $T_{m_n}(\frac{1}{2}\omega_n)$.

3 Discussion

3.1 Regarding assumption (1)

Consider the "line" graphs with $\mathbf{V}_n = \{1, 2, \ldots, n\}$ and $w_{i,i+1} = 2^{-i}$. Here it is not possible to rescale time so that assumption (1) holds. Heuristically, failure of assumption (1) relates to this kind of exponential slowdown of edge-weights at the time of formation of the incipient giant component.

3.2 An example

Let us replace the assumption (1) by the assumption

$$\lim_{t\to 0} \limsup_n \mathbb{E}C_n(t)/|\mathbf{V}_n| = 0; \quad \lim_{t\to\infty} \liminf_n \mathbb{E}C_n(t)/|\mathbf{V}_n| = 1. \tag{13}$$

This is essentially saying, via a compactness argument, that the processes $(C_n(t)/|\mathbf{V}_n|, 0 \le t < \infty)$ converge to a limit process $(\widetilde{C}_\infty(t), 0 \le t < \infty)$ for which $\lim_{t\to 0} \widetilde{C}_\infty(t) = 0$ and $\lim_{t\to\infty} \widetilde{C}_\infty(t) = 1$.

Consider, for fixed $0 < s < 1$, the random times

$$T_n^*(s) := \inf\{t : C_n(t) \ge s|\mathbf{V}_n|\}. \tag{14}$$

Assumption (13) implies these times are $O(1)$ as $n \to \infty$, but is *not* sufficient to show the "weak concentration" property

$$\text{var } T_n^*(s) \to 0 \text{ as } n \to \infty, \tag{15}$$

as the following example shows.

For the complete graph on m vertices with edge-weights $1/m$, our bond percolation process is essentially just the Erdős-Rényi process $\mathcal{G}(m, t/m)$, for which the limit of the process $(C_m(t)/m, 0 \le t < \infty)$ is a certain continuous function $\theta(\cdot)$ with $\theta(t) = 0$ for $t \le 1$ and $\theta(t) > 0$ for $t > 1$ (explicitly, $\theta(t)$ is the solution of $1-\theta = \exp(-t\theta)$ – see e.g. [5] sec. 10.4). Now take two copies of that complete graph on m vertices with edge-weights $1/m$, and add a single edge e^* between them with weight m. This gives a graph on $n = 2m$ vertices. It is easy to see that the limit of the process $(C_n(t)/n, 0 \le t < \infty)$ is the random process

$$\begin{aligned} \widetilde{C}_\infty(t) &= \theta(t)/2, \ 0 \le t < \zeta \\ &= \theta(t), \ \zeta \le t < \infty \end{aligned} \tag{16}$$

where $\mathbb{P}(\zeta \le t) = \theta^2(t)$; here ζ represents the first time at which the giant component in each half contains the end-vertex of e^*. So the "weak concentration" property (15) does not hold.

Note that we can modify this construction to make the graph *vertex-transitive*, that is there is a graph automorphism that maps any vertex to any other vertex. Instead of a single edge between the two original copies, we assign weight $1/m^2$ to *every* such edge. Now the limit is again of form (16) where now $\mathbb{P}(\zeta \le t) = 1 - \exp(-t\theta^2(t))$.

3.3 Analogies with bond percolation on infinite graphs

Rigorous mathematical treatment of bond percolation has focussed on infinite graphs, with "general theory" developed under the assumption of *transitivity*, that is spatial symmetry. As the 2006 survey [15] says,

> ... infinite graphs, where the issue of uniqueness of the giant component translates naturally into the question of whether there is a unique infinite cluster. This has the advantage of always having a clear-cut yes/no answer, in contrast to the finite setting where it is not always totally obvious what one really should mean by a giant component.

In our setting of a *sequence* of finite edge-weighted graphs, one can readily formalize the idea of giant components being unique as the property

$$\sup_t C_n^{[2]}(t)/|\mathbf{V}_n| \to_p 0 \text{ as } n \to \infty \tag{17}$$

where $C_n^{[2]}(t)$ is the size of the *second*-largest component at time t. Note that this can be restated in terms of the jumps of $(C_n(\cdot))$, as

$$\sup_{0 \le t < \infty} |C_n(t) - C_n(t-)|/|\mathbf{V}_n| \to_p 0 \text{ as } n \to \infty. \tag{18}$$

3.4 Two conjectures

Under the background assumption (13), what further assumptions might be sufficient to imply either the "unique giant component" property (17) or the "weak concentration" property (15)? The example at the end of section 3.2 shows that *vertex-transitive* is not sufficient for either property. If instead we assume *edge-transitive* then all edge-weights are equal and so we are in the more familiar setting of bond percolation on an *unweighted* graph with symmetry. Here it seems likely that known methods used in the infinite setting will be relevant. However this requires care: the infinite r-regular tree does not have the "unique giant component" property but typical realizations of random r-regular graphs, as $n \to \infty$, do have this property [18], even though their local weak limit is the infinite r-regular tree. Here is a bold conjecture.

Conjecture 4 *Consider a sequence of edge-transitive graphs with* $|\mathbf{V}_n| \to \infty$. *Then we can always rescale the edge-weight so that (13) holds. After such rescaling, the "weak concentration" and the "unique giant component" properties hold.*

A second bold conjecture is that, without any assumption of symmetry, one of these properties implies the other.

Conjecture 5 *Under assumption (13), the "unique giant component" property (17) implies the "weak concentration" property (15).*

Essentially, the conjecture is saying that the limit process $C_\infty(\cdot)$ indicated at the start of section 3.2 might be deterministic and continuous (as in the classical settings) or might be random and discontinuous (as in the examples in section 3.2), but cannot be random and continuous. (Readers aware of the famous open problems involving continuity of the percolation function on infinite graphs should note that the "weak concentration" property relates to the *inverse* of that function).

Separate from the literature on scaling windows mentioned in section 1, there is a line of work including [4, 7] on bond percolation for unweighted finite graphs under isoperimetry assumptions, that is for expanders, which includes results on uniqueness of giant component. Conjecture 1.1 of [4], not involving isoperimetry assumptions, is somewhat similar to our Conjecture 4.

3.5 Are there analogous results for first passage percolation?

Our starting structure was a finite connected graph $(\mathbf{V}, \mathbf{E})$ with edge-weights $\mathbf{w} = (w_e)$ and with independent Exponential(rate w_e) random variables ξ_e associated with the edges. This structure can alternatively be used to construct first-passage times $X(v, v')$, defined as the minimum of $\sum_{e\in\pi} \xi_e$ over all paths π from v to v'. Regarding this as a model for spread of infection from an initial site v, the set of infected sites at time t is

$$\mathcal{S}(v, t) := \{v' : X(v, v') \le t\}.$$

Write

$$\Delta := \max_{v,v'} \mathbb{E}X(v, v').$$

Given a sequence of such graphs with $|\mathbf{V}_n| \to \infty$ and a sequence $v_n \in \mathbf{V}_n$, consider the times

$$T_n(v_n, s) := \inf\{t : \ |\mathcal{S}_n(v_n, t)| \ge s|\mathbf{V}_n|\}; \quad 0 < s < 1 \tag{19}$$

and the "incipient pandemic" time

$$T_n(v_n) := \inf\{t : \ |\mathcal{S}_n(v_n, t)| \geq |\mathbf{V}_n|/\omega_n\}. \tag{20}$$

There is a simple analysis of such times, provided we impose another assumption. Setting

$$w^* := \min\{w_e \ : \ w_e > 0\}$$

we have

Lemma 6 var $X(v, v') \leq \mathbb{E}X(v, v')/w^*$.

Bounds of this type are classical on $\mathbb{Z}^d$ [16] and are at least folklore in more general settings: an explicit statement and martingale proof in our setting is given in [2]. But inspecting the proof of Proposition 7 in [2] shows that the same bound

$$\text{var } T_n(v_n, s) \leq \mathbb{E}T_n(v_n, s)/w_n^* \tag{21}$$

holds for $T_n(v_n, s)$ at (19), for arbitrary s.

Now assume that in a sequence of graphs

$$w_n^*\Delta_n \to \infty. \tag{22}$$

Because $\mathbb{E}T_n(v_n, s) = O(\Delta_n)$ for fixed s, (21) and (22) imply the "weak concentration" property

$$\frac{T_n(v_n, s)}{\Delta_n} - \frac{\mathbb{E}T_n(v_n, s)}{\Delta_n} \to_p 0.$$

Compactness arguments as in section 2.3 then lead to a conclusion analogous to Proposition 1 for the "incipient pandemic" time:

Corollary 7 *Under assumption (22), there exists a deterministic sequence* $\tau_n(v_n) \in [0, 1]$ *such that, for every sequence* $\omega_n \uparrow \infty$ *sufficiently slowly, the random times* $T_n(v_n)$ *at (20) satisfy*

$$T_n(v_n)/\Delta_n - \tau_n(v_n) \to_p 0.$$

However, this result in the first-passage percolation setting differs in two respects from the Proposition 1 result in the bond percolation setting. To make Corollary 7 interesting we want the sequence $\tau_n(v_n)$ to be bounded away from zero, which is tantamount to the assumption (analogous to the first part of (1)) that for some $0 < t_1 < \infty$

$$\lim_n \mathbb{E}|\mathcal{S}_n(v_n, t_1\Delta)|/|\mathbf{V}_n| = 0. \tag{23}$$

But in classical settings such as nearest-neighbor first-passage percolation on $\mathbb{Z}_m^d$ [17, 6] this does not hold, because by the shape theorem the scaling limit of

$|\mathcal{S}_n(v_n, t\Delta)|/|\mathbf{V}_n|$ is a deterministic function $\phi(t)$ with $\phi(t) > 0$ for $t > 0$. The context where we do expect (23) to hold is where the epidemic starts with faster than polynomial growth, for instance on expander graphs or familiar models of random graphs [8, 9]. Second, while assumption (22) is stronger than necessary, we do need some assumption to prevent $T_n(v_n)$ have variability due to the influence of a single edge-traversal time ξ_e associated with a very small weight w_e. For weak concentration of point-to-point percolation times $X(v, v')$, precise conditions in terms of such influence are given in [2]. It seems plausible that there are analogous precise conditions for weak concentration of the "incipient pandemic" time $T_n(v_n)$, but we have not studied this issue.

Acknowledgements. I thank Asaf Nachmias for helpful comments.

References

[1] David Aldous. Interacting particle systems as stochastic social dynamics. *Bernoulli*, 19(4):1122–1149, 2013.

[2] D.J. Aldous. Weak concentration for first passage percolation times on graphs and general increasing set-valued processes. arXiv e-print 1604.06418, 2016.

[3] D.J. Aldous and Xiang Li. A framework for imperfectly observed networks, 2016. In Preparation.

[4] Noga Alon, Itai Benjamini, and Alan Stacey. Percolation on finite graphs and isoperimetric inequalities. *Ann. Probab.*, 32(3A):1727–1745, 2004.

[5] Noga Alon and Joel H. Spencer. *The probabilistic method*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons, Inc., Hoboken, NJ, third edition, 2008. With an appendix on the life and work of Paul Erdős.

[6] Antonio Auffinger, Jack Hanson, and Michael Damron. 50 years of first passage percolation, 2015. eprint arXiv:1511.03262.

[7] Itai Benjamini, Stéphane Boucheron, Gábor Lugosi, and Raphaël Rossignol. Sharp threshold for percolation on expanders. *Ann. Probab.*, 40(1):130–145, 2012.

[8] Shankar Bhamidi, Remco van der Hofstad, and Gerard Hooghiemstra. First passage percolation on random graphs with finite mean degrees. *Ann. Appl. Probab.*, 20(5):1907–1965, 2010.

[9] Shankar Bhamidi, Remco van der Hofstad, and Gerard Hooghiemstra. First passage percolation on the Erdős-Rényi random graph. *Combin. Probab. Comput.*, 20(5):683–707, 2011.

[10] Béla Bollobás and Oliver Riordan. The phase transition in the Erdős-Rényi random graph process. In *Erdös centennial*, volume 25 of *Bolyai Soc. Math. Stud.*, pages 59–110. János Bolyai Math. Soc., Budapest, 2013.

[11] C. Borgs, J. T. Chayes, H. Kesten, and J. Spencer. The birth of the infinite cluster: finite-size scaling in percolation. *Comm. Math. Phys.*, 224(1):153–204, 2001. Dedicated to Joel L. Lebowitz.

[12] Christian Borgs, Jennifer T. Chayes, Remco van der Hofstad, Gordon Slade, and Joel Spencer. Random subgraphs of finite graphs. I. The scaling window under the triangle condition. *Random Structures Algorithms*, 27(2):137–184, 2005.

[13] Christian Borgs, Jennifer T. Chayes, Remco van der Hofstad, Gordon Slade, and Joel Spencer. Random subgraphs of finite graphs. II. The lace expansion and the triangle condition. *Ann. Probab.*, 33(5):1886–1944, 2005.

[14] Fan Chung, Paul Horn, and Linyuan Lu. Percolation in general graphs. *Internet Math.*, 6(3):331–347 (2010), 2009.

[15] Olle Häggström and Johan Jonasson. Uniqueness and non-uniqueness in percolation theory. *Probab. Surv.*, 3:289–344, 2006.

[16] Harry Kesten. On the speed of convergence in first-passage percolation. *Ann. Appl. Probab.*, 3(2):296–338, 1993.

[17] Harry Kesten. First-passage percolation. In *From Classical to Modern Probability*, volume 54 of *Progr. Probab.*, pages 93–143. Birkhäuser, Basel, 2003.

[18] Asaf Nachmias and Yuval Peres. Critical percolation on random regular graphs. *Random Structures Algorithms*, 36(2):111–148, 2010.